

Beyond Counting: New Perspectives on the Active IPv4 Address Space

@IETF 96 Berlin (*maprg*)
July 2016

Philipp Richter
TU Berlin

Georgios Smaragdakis
MIT

David Plonka
Akamai


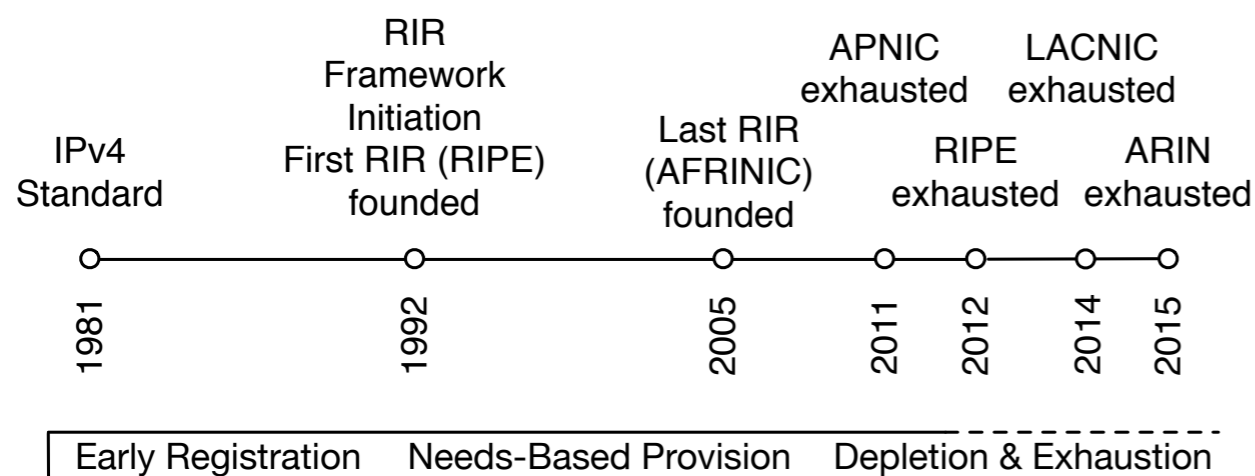
Arthur Berger
Akamai/MIT

*work under submission
comments highly appreciated!*

preprint: <http://arxiv.org/abs/1606.00360>

IPv4 Address Space Exhaustion

- IPv4 has been around for ~35 years
- Theoretically routable IP addresses: 3.7B, ~2.8B routed
- IANA exhausted its address pool in 2011
- Today: Less than 2% of the IPv4 address space “free”



Figures: P. Richter, M. Allman, R. Bush, V. Paxson: A Primer on IPv4 Scarcity, *ACM CCR* 45(2), 2015.

Operators' Community Efforts



Efforts in the IETF community:

- IPv6 transition mechanisms
- IPv4 multiplexing/sharing mechanisms (e.g., EnIP, A+P)
- Efforts to conserve IPv4 address space

e.g., [draft-fleischhauer-ipv4-addr-saving-05](#), RFC6346, [draft-chimiak-enhanced-ipv4-03](#)

IANA/Regional Internet Registries:

- Establishment of address transfer policies
- Incentives for increasing address space utilization

Academic Community Efforts



- Measurements to understand “where we are” right now
- Internet-wide: Number of actively used IPv4 addresses:
 - “1.2B IP addresses in use in 2014”, statistical estimation Zander et al., IMC ‘14
 - “5.3M /24 address blocks in use in 2013”, passive+active measurement Dainotti et al., JSAC ‘16
- Challenge: No single vantage point captures all activity

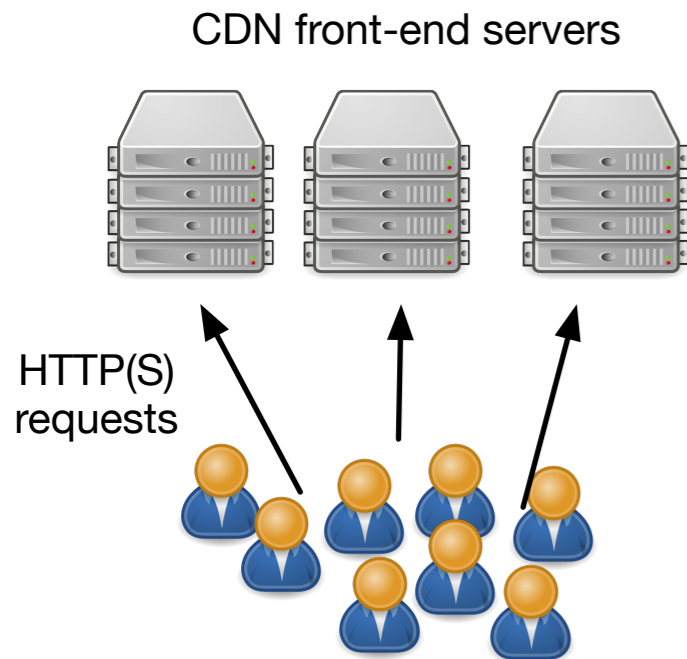
Academic Community Efforts



- Measurements to understand “where we are” right now
- Internet-wide: Number of actively used IPv4 addresses:
 - “1.2B IP addresses in use in 2014”, statistical estimation Zander et al., IMC ‘14
 - “5.3M /24 address blocks in use in 2013”, passive+active measurement Dainotti et al., JSAC ‘16
- Challenge: No single vantage point captures all activity

**What can we say from our CDN’s perspective?
Can we do more than counting active IP addresses?**

The CDN as an Observatory

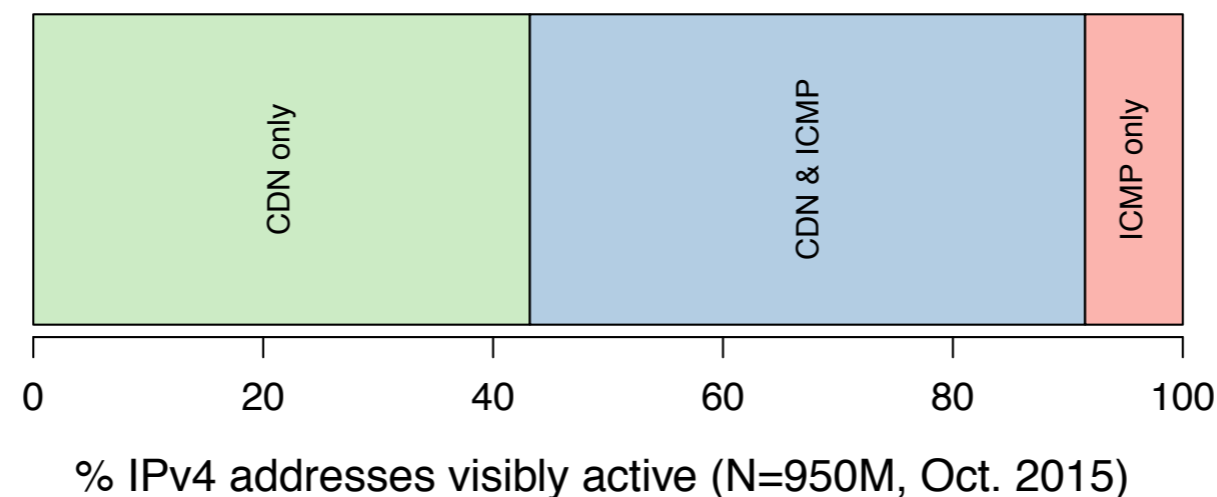


- 200,000+ servers
- 3 trillion requests per day
- CDN logs: number of requests per IP per day


Totals for the entirety of 2015:

- 1.2B active IPv4 addresses (42% of routed)
- 6.5M active /24 address blocks (59% of routed)

Visibility: CDN logs vs. ICMP scan (ZMap project, 8 snapshots)

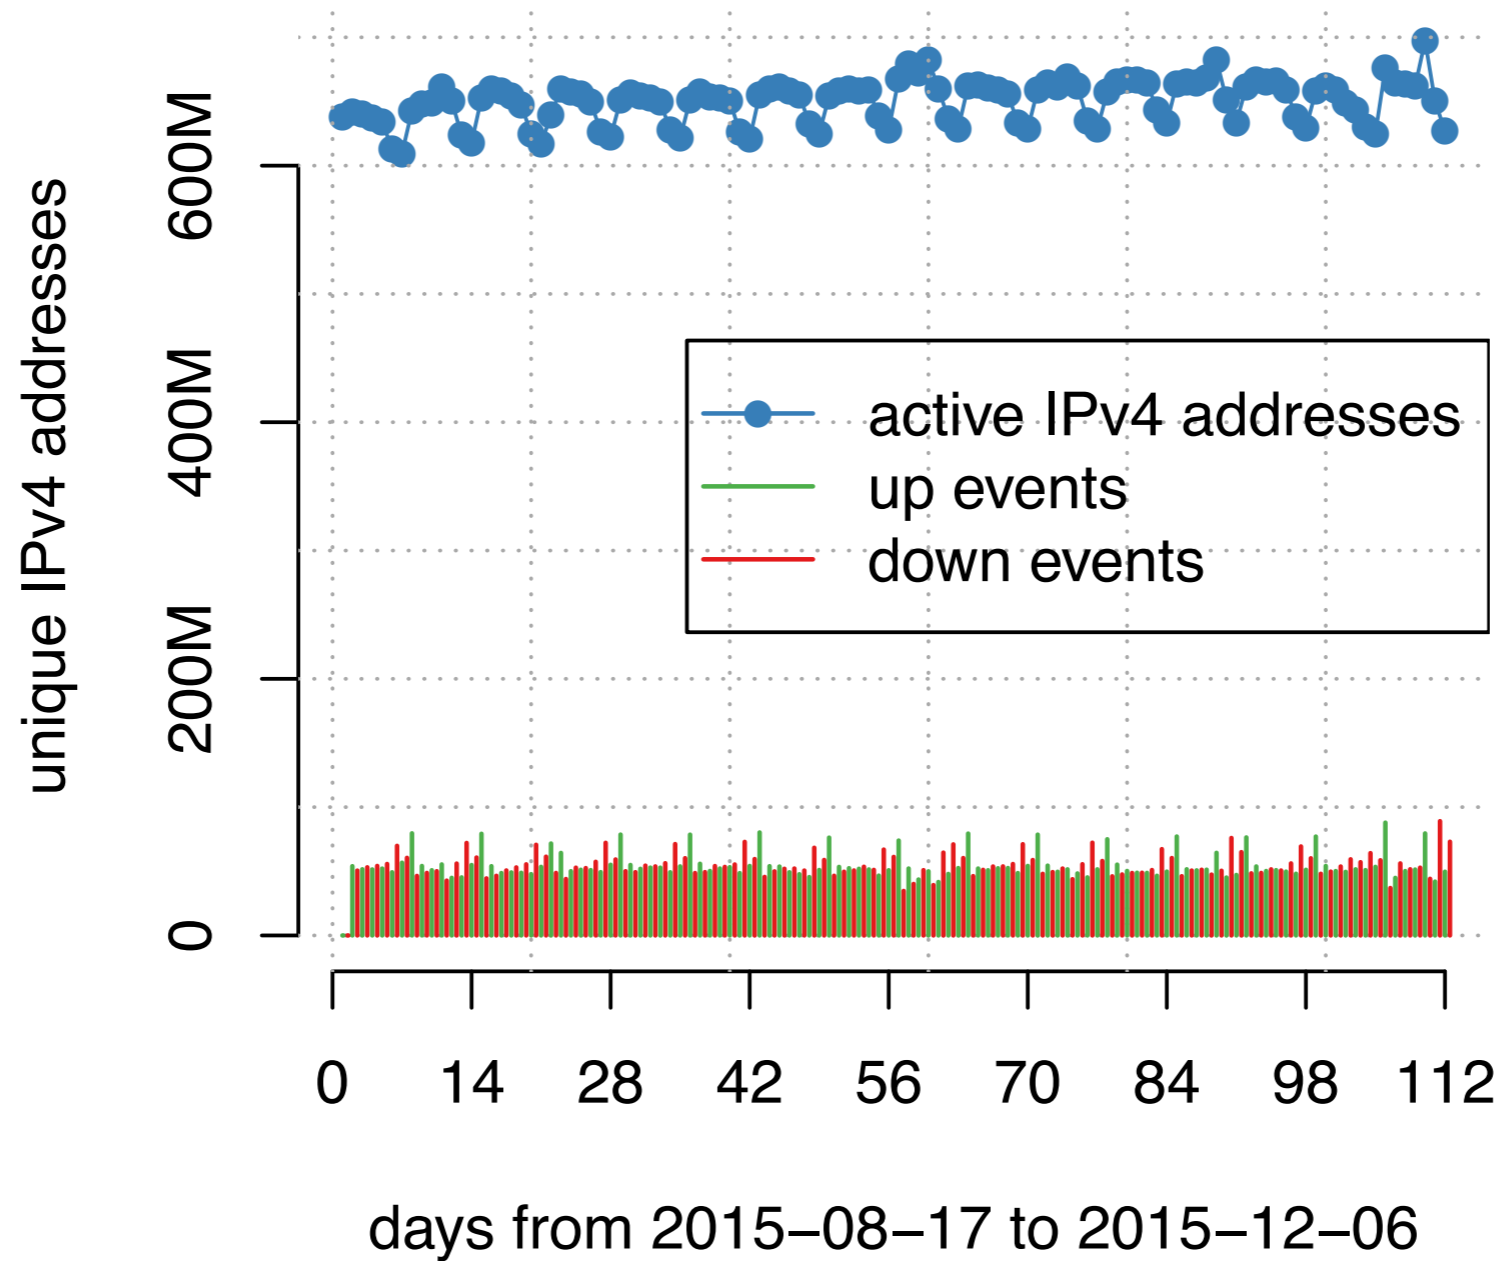


Peak IPv4?



Active IPv4 address counts have stagnated since 2014

Daily IPv4 Activity and Churn



Churn on all Timescales

day-to-day: ~7% come, ~7% go



week-to-week: ~5% come, ~5% go




month-to-month: ~5% come, ~5% go




**The *number* of active IPv4 addresses stays constant
the *set* of active addresses varies on all timescales**

Long-term Effect of Address Churn




Over the course of one year, 25% of the active IP address pool changed

Address Activity Matrix




for each day on which an IP address was active (requested content), we draw a red dot.


Patterns: “In situ” Address Activity




static block DE University



DHCP pool US University



residential users US ISP




residential users DE ISP


**“in situ” activity:
address assignment practice
+
user behavior**

**(no visible modification of
address assignment practice)**

Patterns: Operational Change



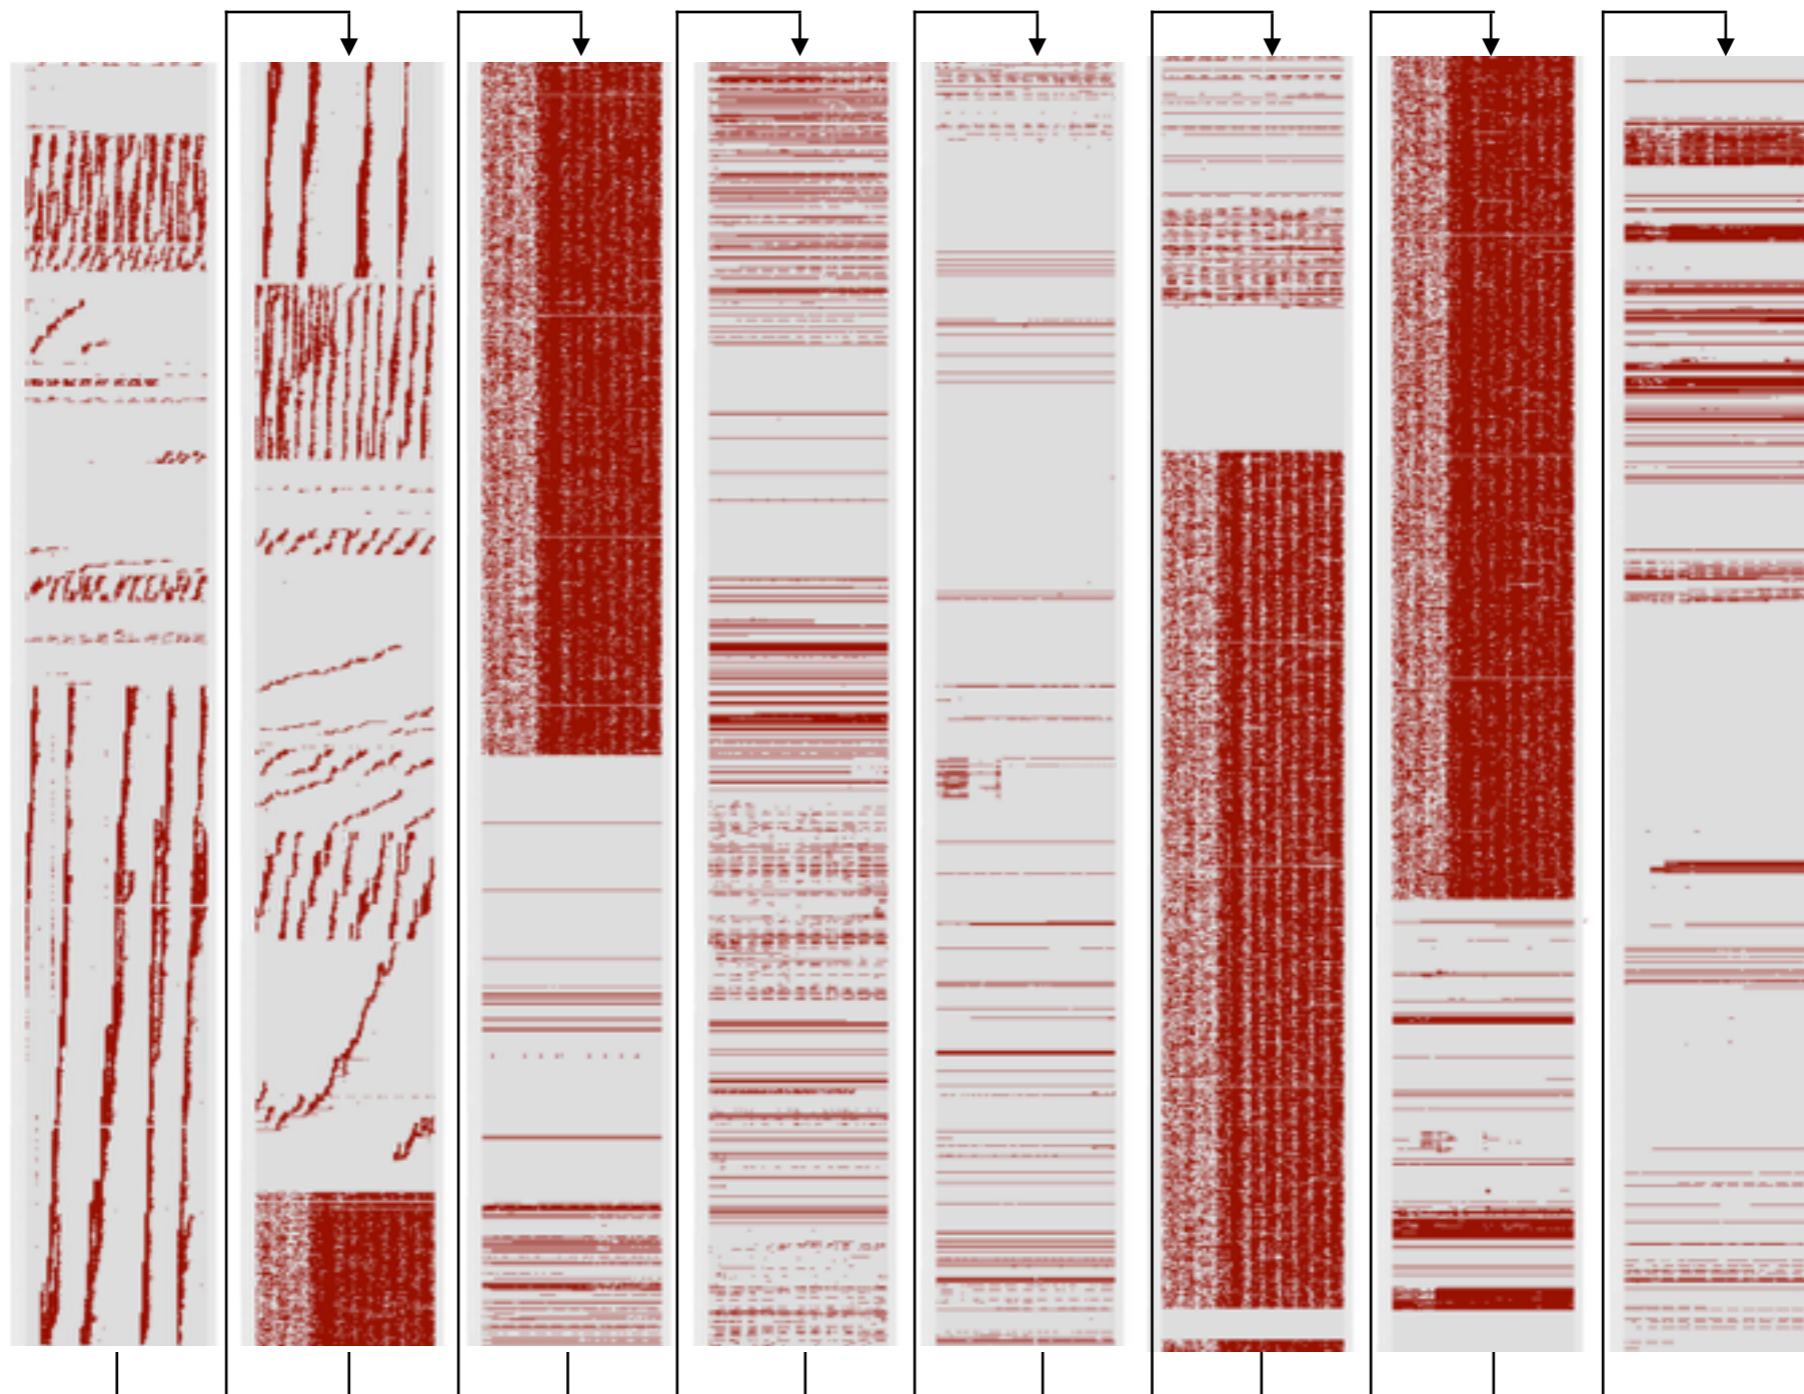
DE University



DE University


Activity Matrix at Scale

20k adjacent IP addresses (in active /24s), University Network




Metric 1: Filling Degree per /24

Number of active IP addresses per /24 [1...256]

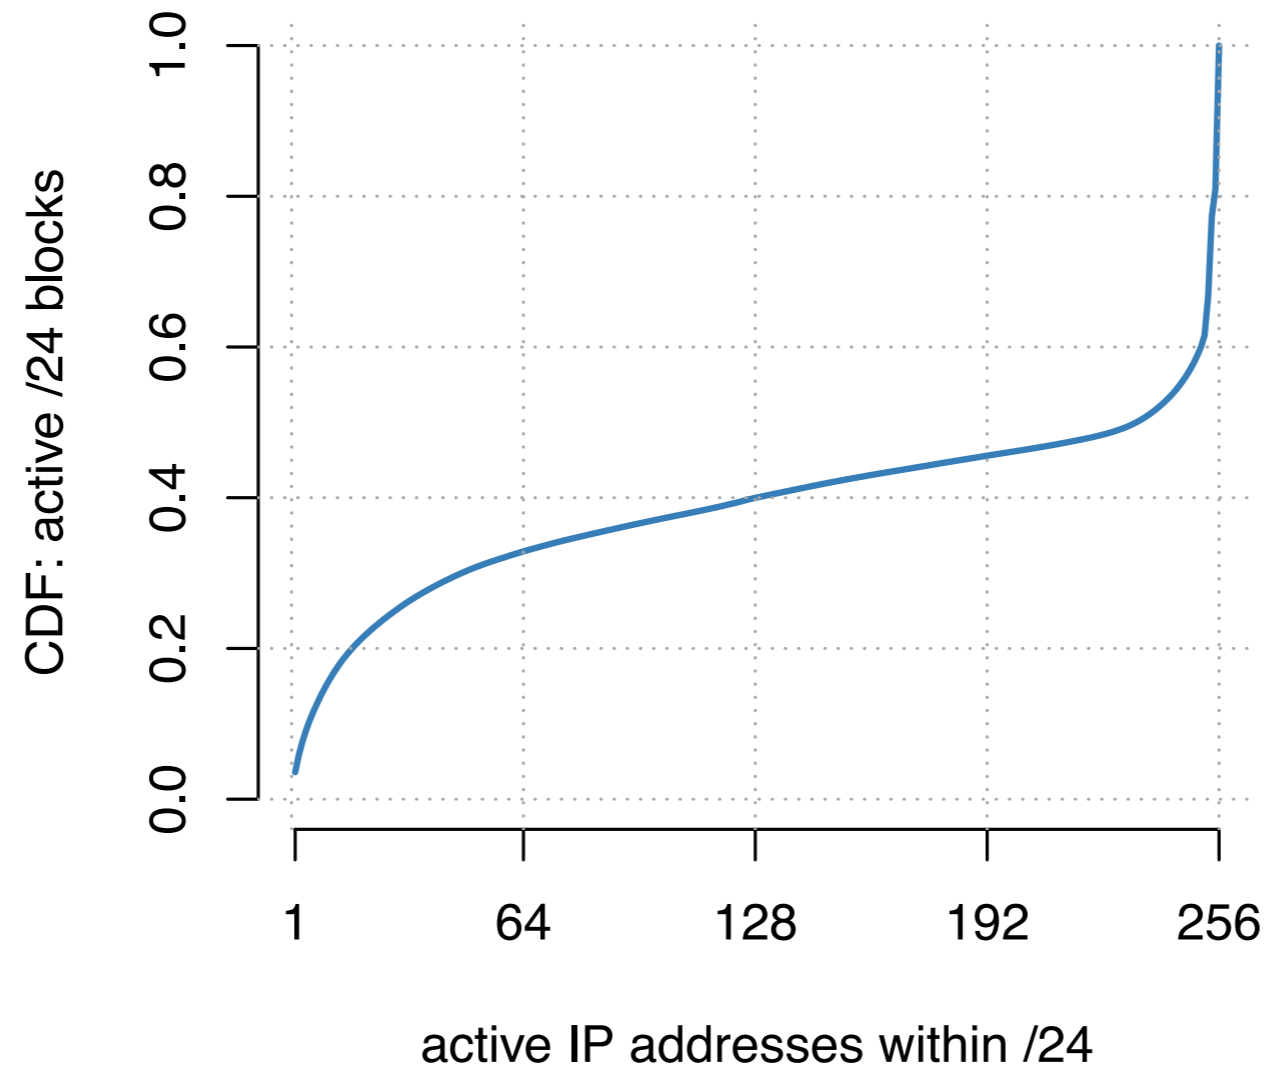


rather low
(degree = 29)

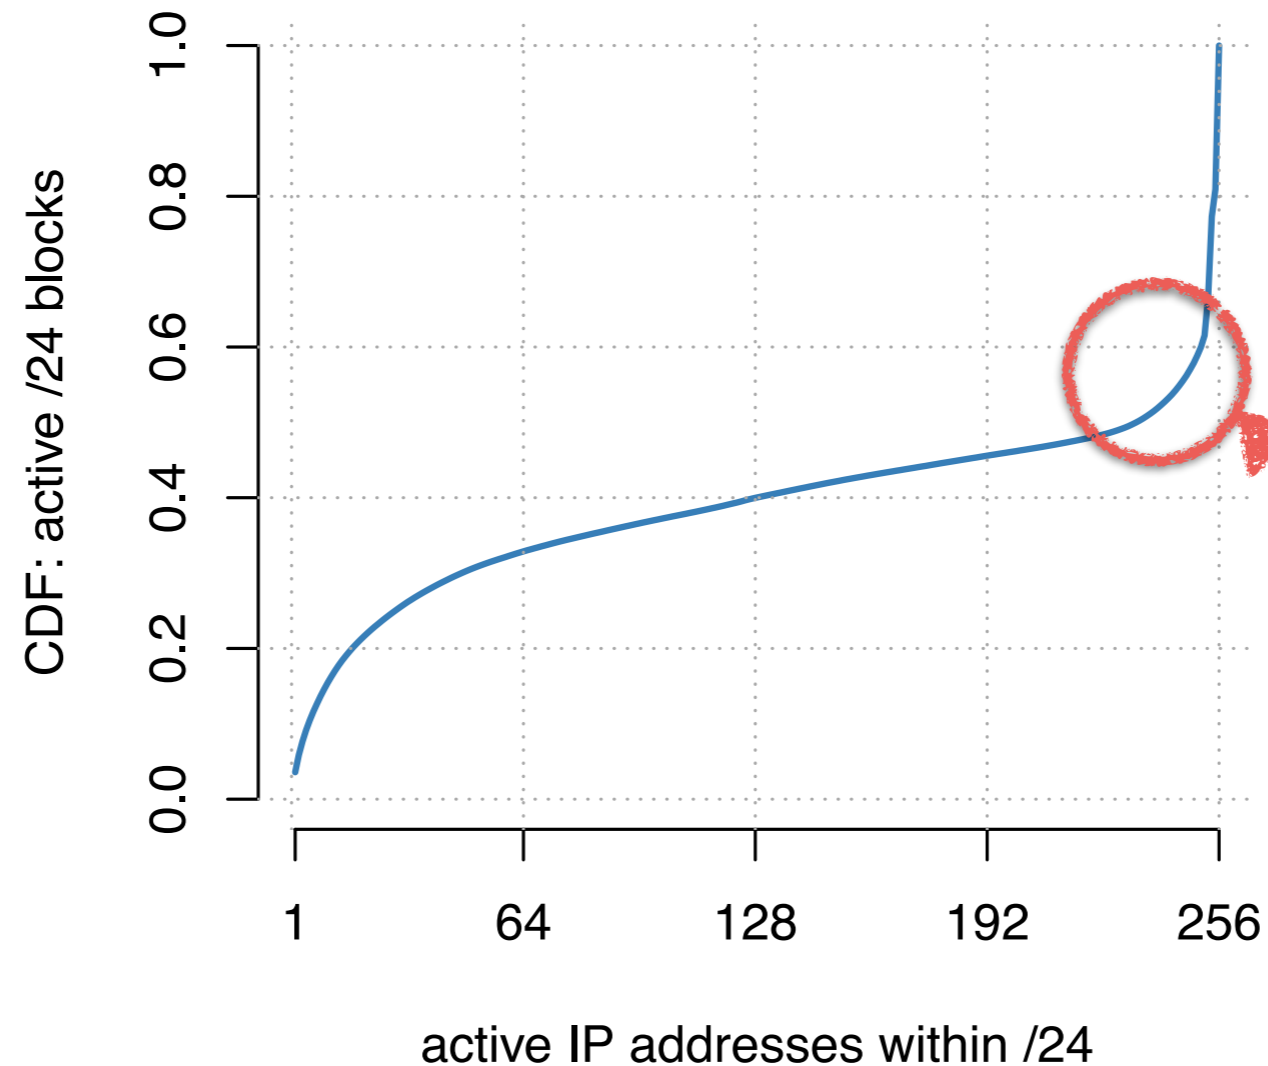


high
(degree = 254)

Metric 1: Filling Degree per /24




Metric 1: Filling Degree per /24



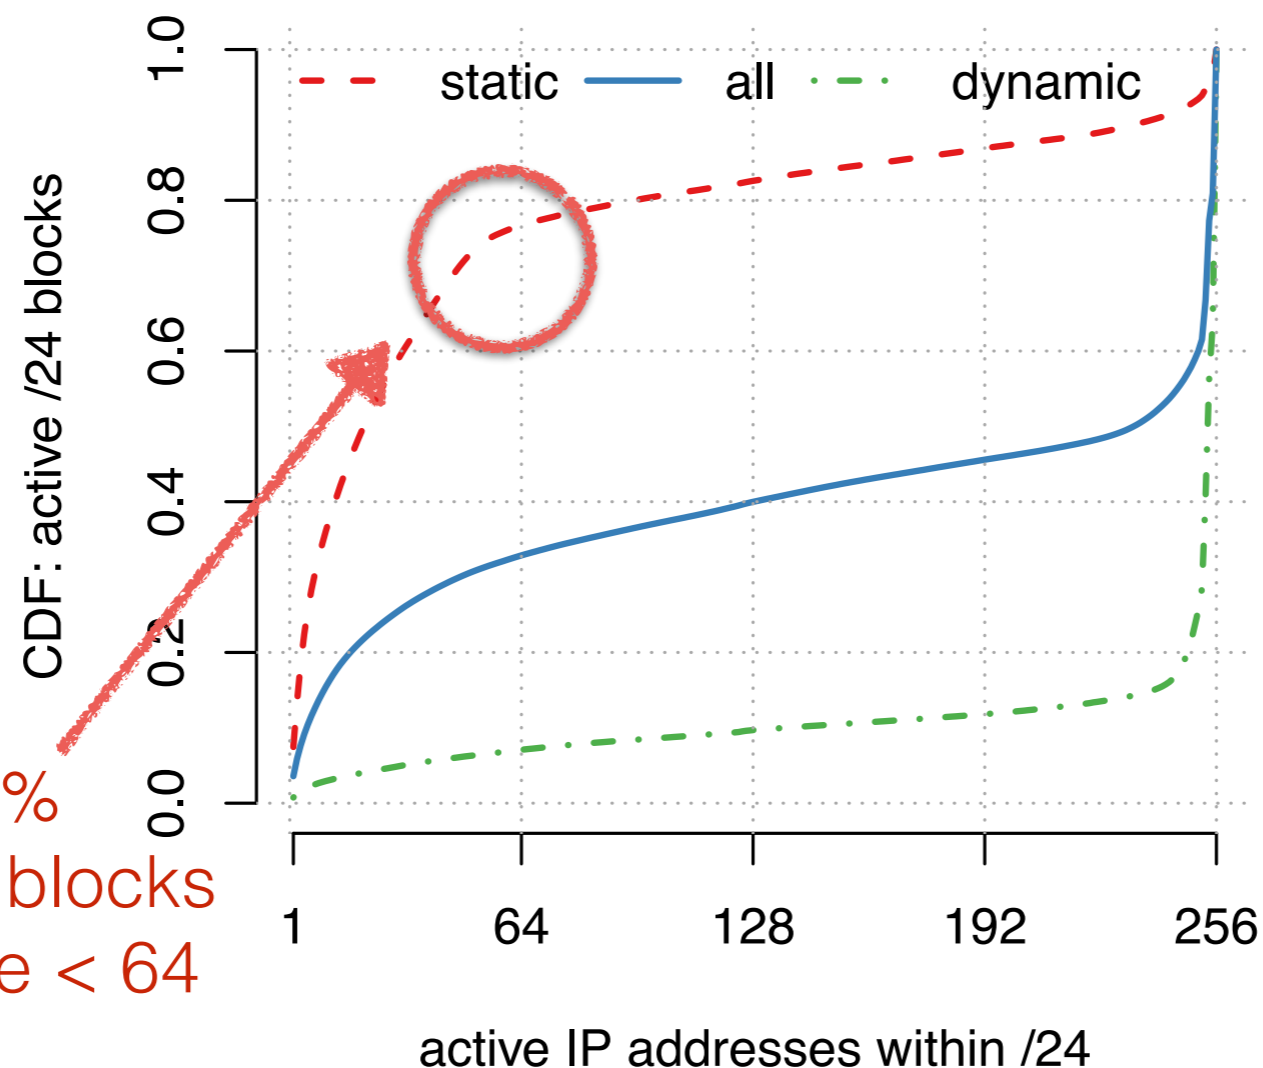
only less than 50% of **all** active /24 blocks have filling degree > 250

Addressing: Static vs. Dynamic



- We tagged likely static/dynamic blocks using PTR records
- We identified 262K static blocks and 456K dynamic blocks


Addressing: Static vs. Dynamic



more than 70%
of "static"-tagged blocks
have filling degree < 64

- We tagged likely static/dynamic blocks using PTR records
- We identified 262K static blocks and 456K dynamic blocks


Addressing: Static vs. Dynamic




- We tagged likely static/dynamic blocks using PTR records
- We identified 262K static blocks and 456K dynamic blocks

Metric 2: Spatio-temporal Utilization

$$\frac{\text{sum}(\langle \text{active IP, day} \rangle)}{\text{sum}(\text{all possible } \langle \text{active IP, day} \rangle)} = \frac{\text{red}}{\text{red} + \text{grey}}$$



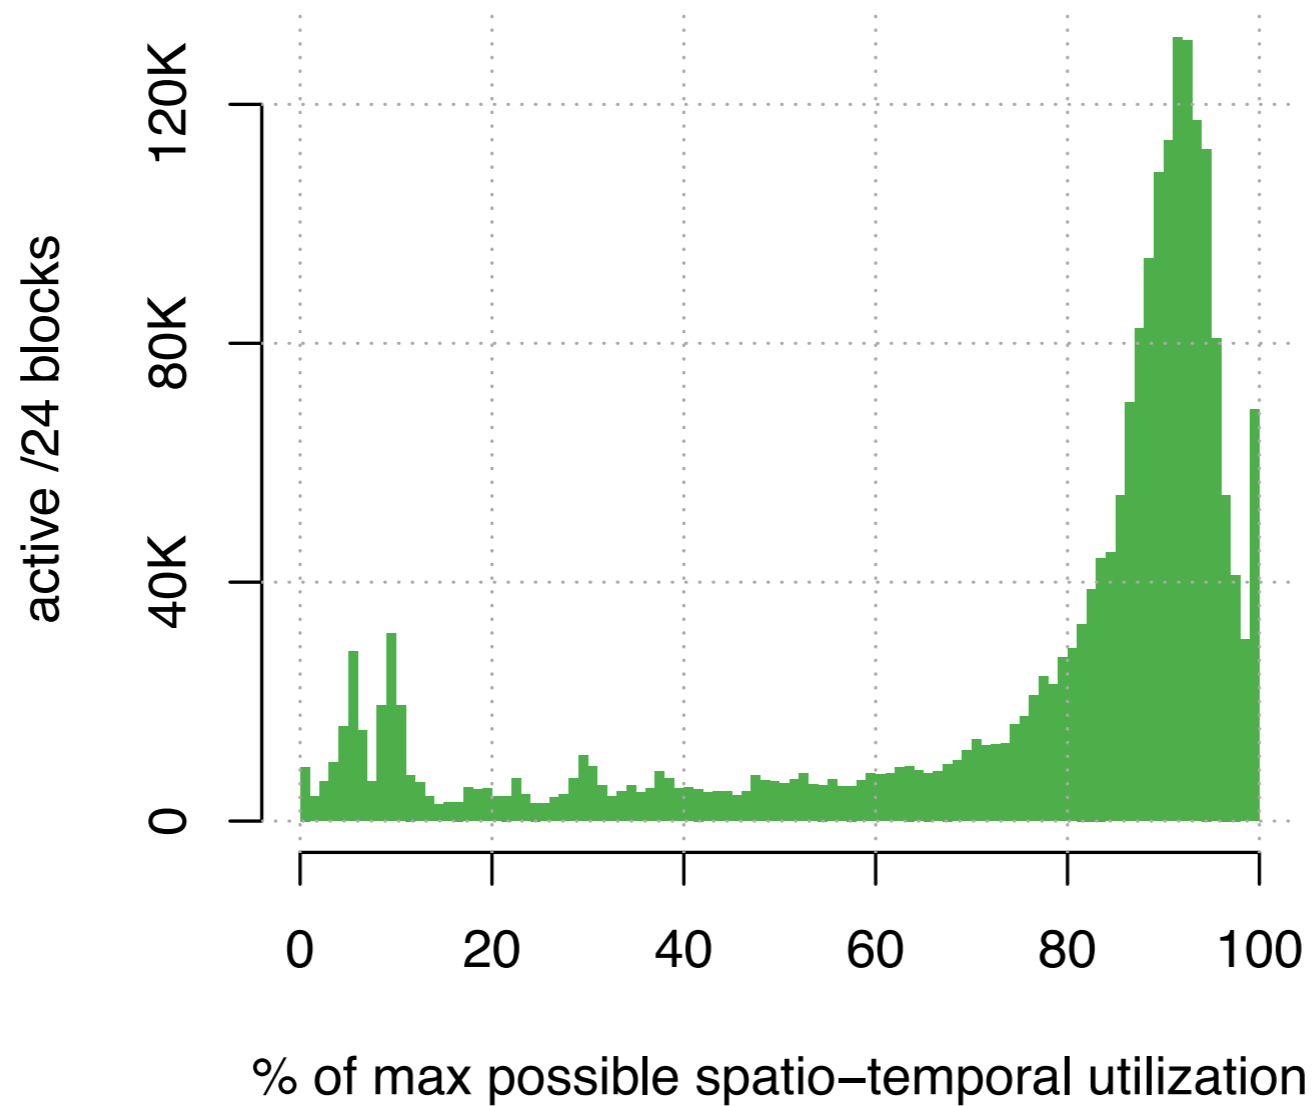
low utilization (18%)



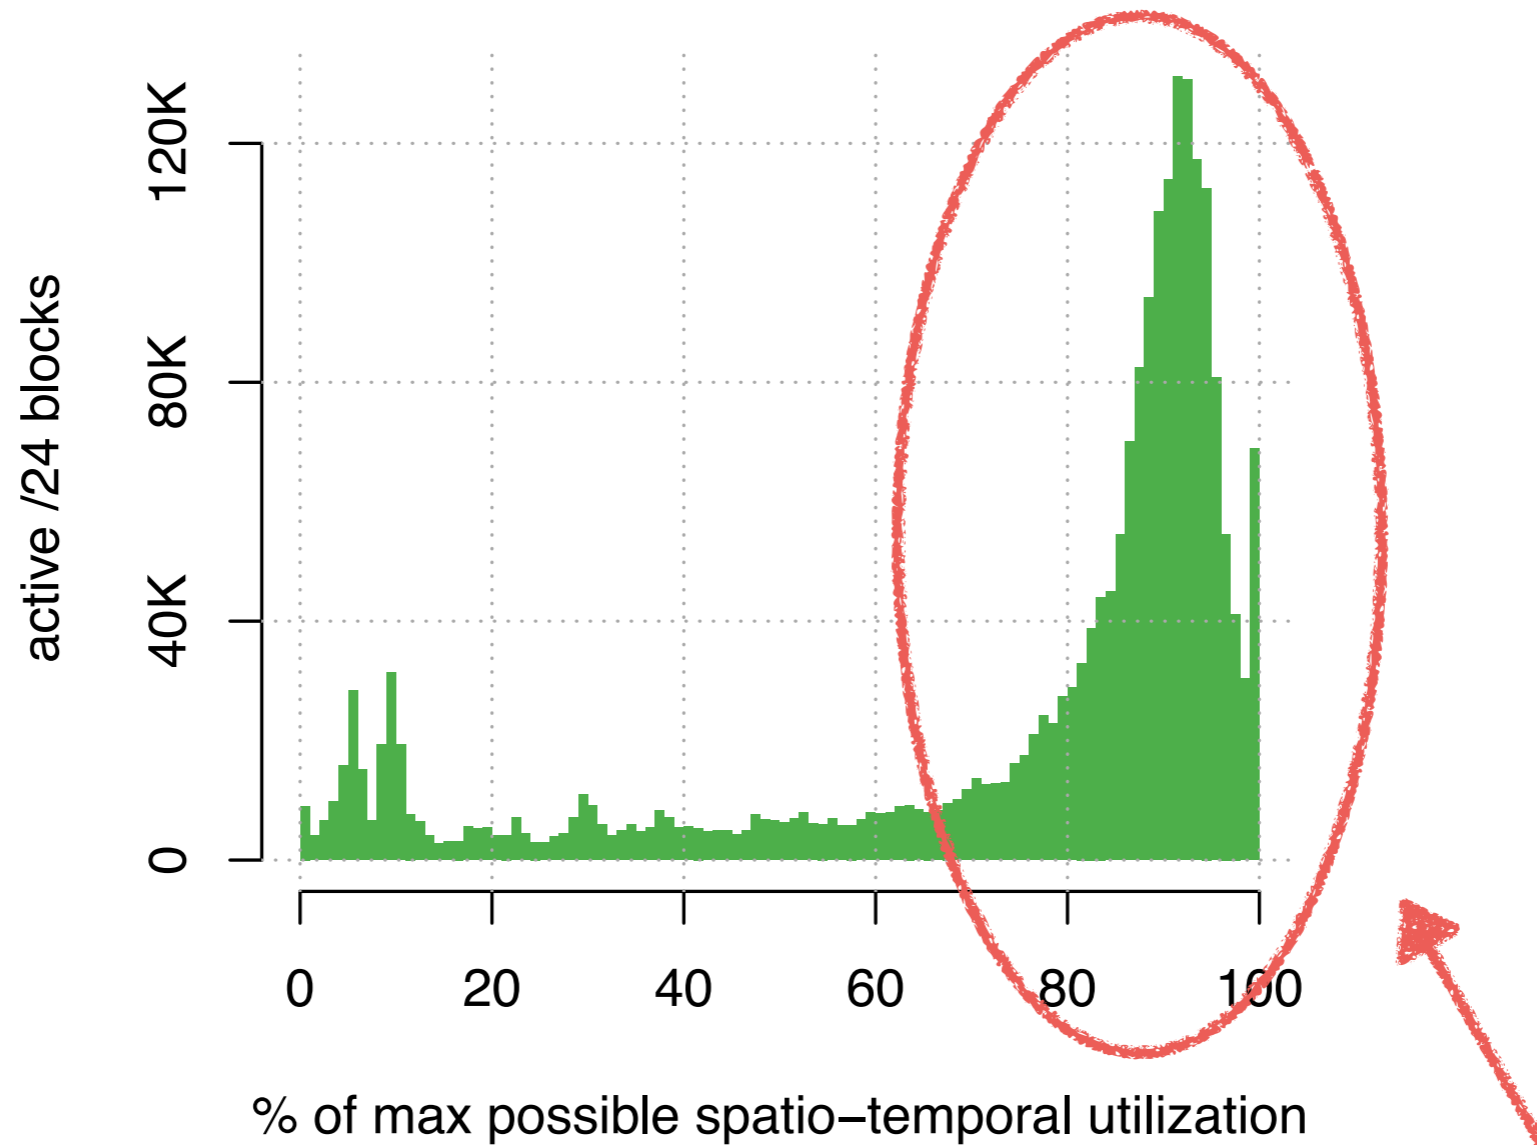
rather high (75%)

Dynamic addressing: Configuration/Pool sizes matter

Utilization: Blocks w/ > 250 active IPs




Utilization: Blocks w/ > 250 active IPs



majority of - likely dynamic - blocks show high utilization

Utilization: Blocks w/ > 250 active IPs




a third of - likely dynamic - blocks show low utilization

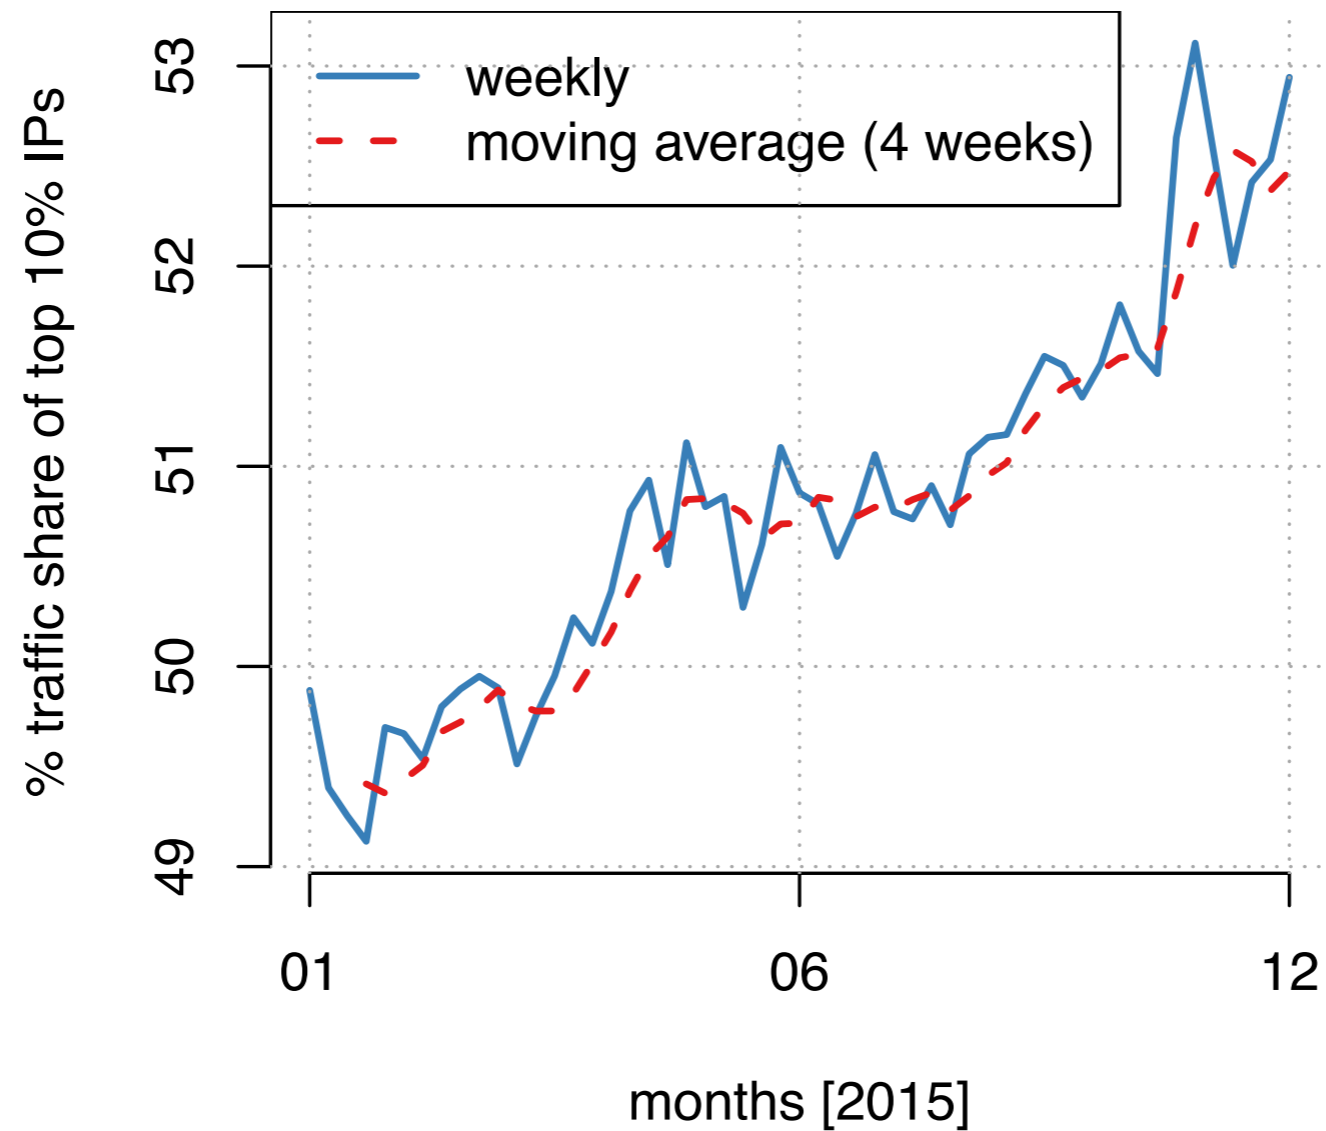
- Comprehensive study of IPv4 address activity
- Metrics “beyond” binary notion of IPv4 activity
- Can inform: Network operations, address [re]assignment
- Can inform: Network security and host reputation

Figure: active /24 address blocks

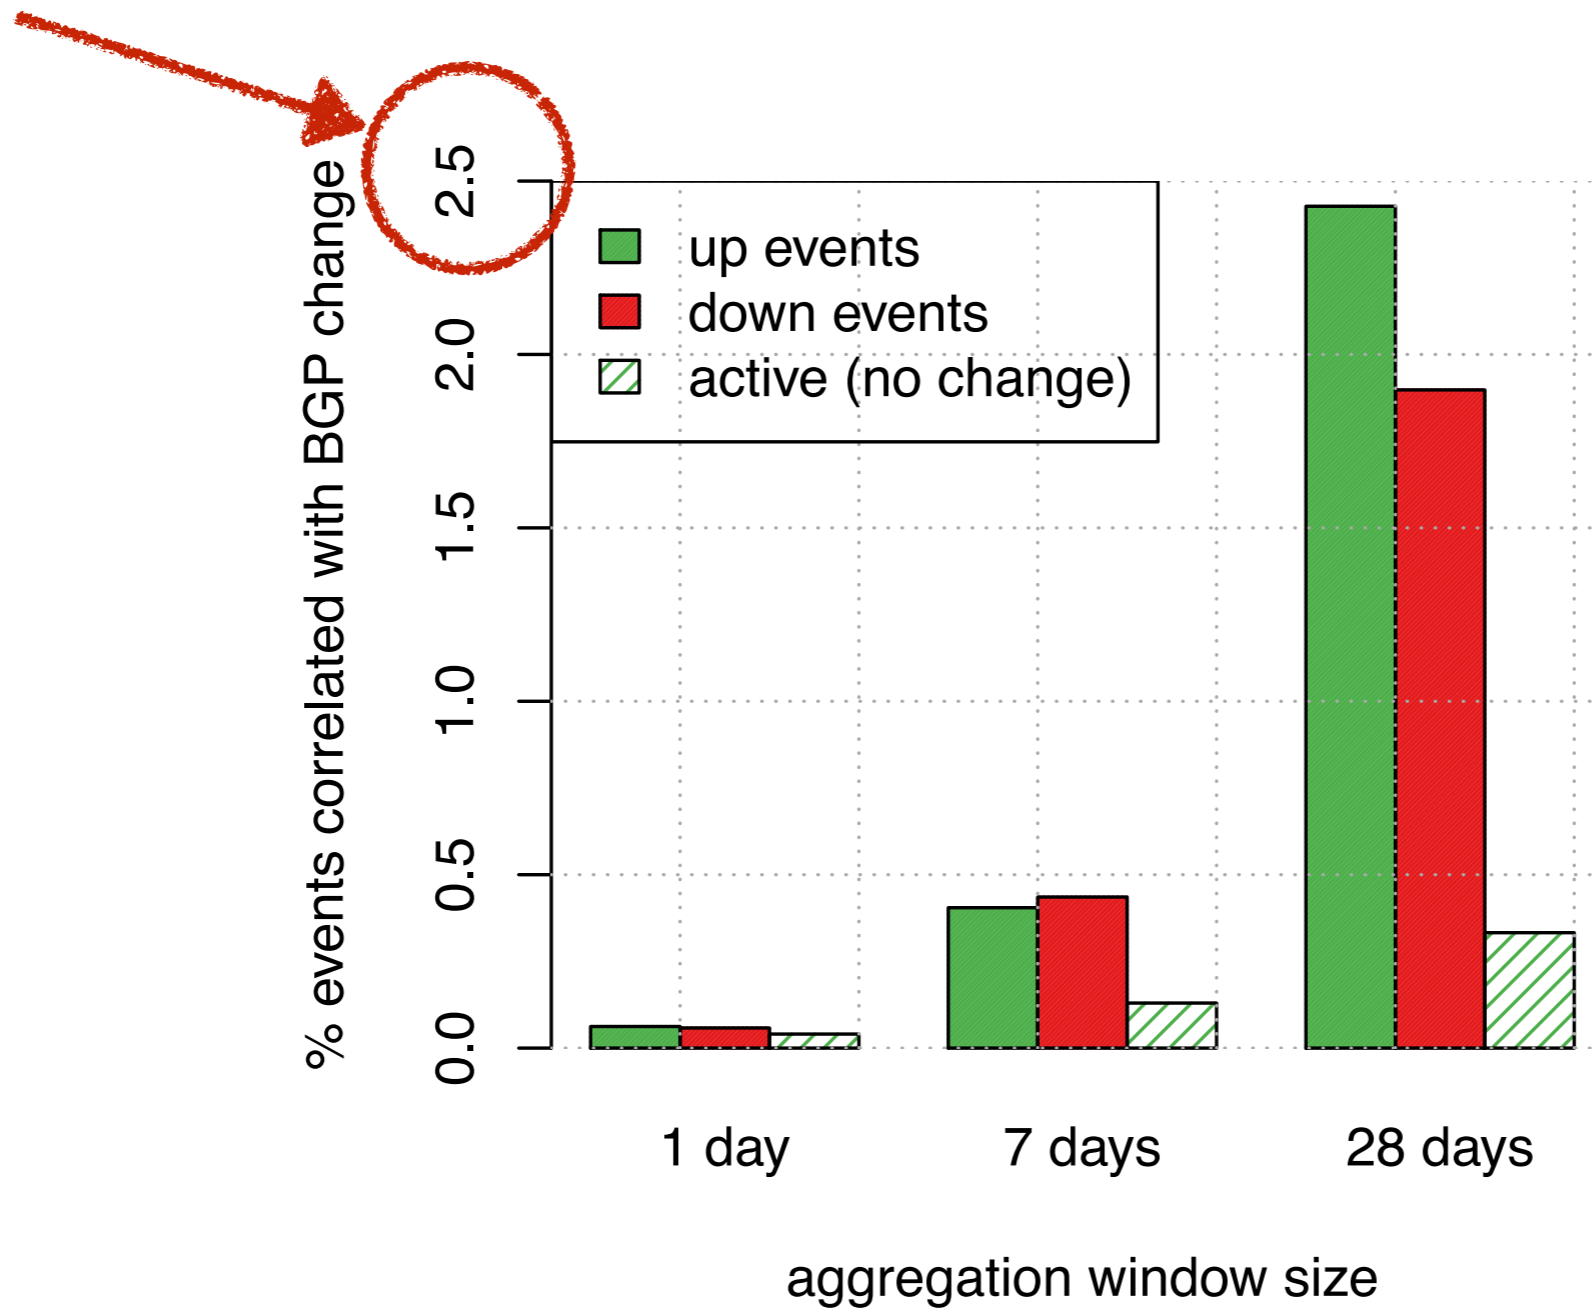
- Spatio-temporal utilization
- Traffic contribution
- Relative host count



Backup: IPv4 Traffic Consolidation




Backup: Churn Visibility in BGP




Backup: Classification ICMP-only IPs



Visibility CDN/ICMP



ICMP-only hosts



server identification: ZMap scans HTTP(S), POP3(S), IMAP(S)

router identification: Ark, TTL exceeded received

Backup: IPv6 /64 Growth

